Optimal Ancilla-free Pauli+V Circuits for Axial Rotations

Andreas Blass,[1] Alex Bocharov,[2] and Yuri Gurevich[2]

[1)] *Mathematics, University of Michigan, Ann Arbor, MI, USA*

[2)] *Microsoft Research, Redmond, WA, USA*

arXiv:1412.1033v1 [quant-ph] 2 Dec 2014

Recently Neil Ross and Peter Selinger analyzed the problem of approximating z-rotations by means of single-qubit Clifford+T circuits. Their main contribution is a deterministic-search technique which allowed them to make approximating circuits shallower.

We adapt the deterministic-search technique to the case of Pauli+V circuits and prove similar results. Because of the relative simplicity of the Pauli+V framework, we use much simpler geometric methods.

I. INTRODUCTION

For any universal basis $\mathcal{B}$ for single-qubit circuits, this natural problem arises: Given a single-qubit gate G and real $\varepsilon > 0$, construct an ancilla-free $\mathcal{B}$-circuit that approximates G with precision ε. There is a well-known elementary reduction of this problem to its special case where G is a z-rotation $R_z(\theta) = \begin{pmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{pmatrix}$ [Ref. 15]. The reduction does not work for all bases but it works for Clifford+T, for Pauli+V and for most other universal single-qubit bases in the literature. We restrict attention to the special case.

The matrix of the approximating $\mathcal{B}$-circuit has the form $\begin{pmatrix} u & -v^* \\ v & u^* \end{pmatrix}$. In many cases, including those of Clifford+T and Pauli+V, the circuit can be efficiently constructed from the matrix. The problem becomes just to find an appropriate pair (u, v) of complex numbers.

In article [Ref. 22], Peter Selinger introduced a randomized-search technique for finding a desired pair (u, v) in the Clifford+T framework. The result was an efficient probabilistic circuit-synthesis algorithm for the Clifford+T basis. "Under a mild hypothesis on the distribution of primes, the expected running time of the probabilistic algorithm is polynomial in $\log(1/\varepsilon)$, and the depth of the resulting approximating circuit is $O(1/\varepsilon)$. If the gate to be approximated is a z-rotation, the T-count of the approximating circuit is $4\log_2(1/\varepsilon)+O(1)$."

Let V be the set of 3 unitary operators

$$V_1 = (I + 2iX)/\sqrt{5}, \quad V_2 = (I + 2iY)/\sqrt{5}, \quad V_3 = (I + 2iZ)/\sqrt{5}$$

where I is the identity matrix and X, Y, Z are the single-qubit Pauli matrices. The group generated by the three operators was introduced and studied in [Refs. 11 and 12]. The use of the V basis for quantum computing was initiated and studied in [Ref. 6].

Using the randomized-search technique, two of the present authors and Krysta Svore developed a probabilistic algorithm analogous to Selinger's for Pauli+V (and thus for Clifford+V) circuits [Ref. 2]. Under a conjecture on the distribution of primes, the expected running time of their algorithm is polynomial in $\log(1/\varepsilon)$, and the depth of the resulting circuit is $4\log_5(1/\varepsilon) + O(1)$. The conjecture is rather credible and purely number-theoretic.

Later, also in the framework of Clifford+T, Ross and Selinger replaced the randomized-search technique with an even more efficient deterministic-search technique [Ref. 20]. Under a hypothesis on the distribution of primes, the expected running time of the new circuit-synthesis algorithm is polynomial in $\log(1/\varepsilon)$, and the T-count of the approximating circuit is

$3\log_2(1/\varepsilon) + O(\log\log(1/\varepsilon))$. If an oracle to factor integers is available (e.g. Shor's factoring algorithm), the approximating circuit has the minimal possible depth.

Ross and Selinger suggested that some other universal quantum bases may be amenable to a similarly optimal synthesis. We show in this paper that this suggestion is realized in the case of the Pauli+V basis. The deterministic-search technique simplifies quite substantially in the Pauli+V case.

We present two circuit-synthesis algorithms using credible number-theoretic conjectures. The first synthesis algorithm runs in expected polynomial time with respect to $\log\frac{1}{\varepsilon}$ and constructs a Pauli+V circuit of depth at most $3\log_5\frac{1}{\varepsilon} + O(\log\log\frac{1}{\varepsilon})$ that approximates the axial rotation $R_z(\theta)$ within ε. The second synthesis algorithm takes an additional input parameter, namely an error tolerance $\delta > 0$. It runs in time polynomial in $\log\frac{1}{\varepsilon}$ and $\log\frac{1}{\delta}$, and it is in fact fast and practical. It returns Nil or constructs a Pauli+V circuit of depth at most $3\log_5\frac{1}{\varepsilon} + O(\log\log\frac{1}{\varepsilon})$ that approximates the axial rotation $R_z(\theta)$ within ε. The probability of returning Nil is at most δ.

It may be important to find a similar solution in the framework of Fibonacci circuits described e.g. in [Ref. 9].

Related Work After our result was pre-announced in [Ref. 3], Neil Ross published another confirmation that the deterministic-search technique works for the Pauli+V basis [Ref. 19]. Our solution is simpler, conceptually and algorithmically.

II. PRELIMINARIES

A. Geometry

Consider the complex plane $\mathbb{C}$. The real and imaginary axes meet at point 0 which is the origin of the coordinate system and will be denoted O. Every nonzero point $z \in \mathbb{C}$ can be viewed as a vector from the origin O to point z. It will be clear from the context when a point is treated as a vector.

If u, v are distinct points of the plane then $[u, v]$ is the straight-line segment between u and v, and $|u, v|$ is the length of $[u, v]$. If, in addition, points u, v are on the unit circle around O and are not the opposites of each other, let $\mathrm{Arc}(u, v)$ be the shorter arc of the unit circle between u and v.

Given a positive real $\varepsilon < 1$, consider a circular segment, or meniscus, $M_0 = \{u : 1-\varepsilon^2 < \Re(u) \leq 1\}$ of the unit disk, centered around point 1. Given a real number θ, rotate M_0 to the angle $-\theta/2$ around O; the resulting meniscus is centered around the point $e^{-i\theta/2}$ and will be denoted $M_\varepsilon(\theta)$ so that $M_0 = M_\varepsilon(0)$. Menisci play important role in our approximation problem.

Given a meniscus $M = M_\varepsilon(\theta)$, centered around the point $z = e^{-i\,\theta/2}$, let z_1, z_2 be the two corner points of the meniscus. Let $z_0 = (z_1 + z_2)/2$ and let z_3 be the intersection point of the tangent lines to the unit circle at z_1 and z_2. The following terminology will be useful.

- The chord $[z_1, z_2]$ is the *base* of M, and the vectors $z_1 - z_2$ and $z_2 - z_1$ are the *base vectors* of M.
- $\mathrm{Arc}(z_1, z_2)$ is the *arc* of M.
- The vector $z - z_0$ is the *handle* of M. Note that the handle uniquely defines the meniscus.
- The isosceles triangle formed by points z_1, z_2, z_3 is the *enclosing triangle* of M. The base of M is also the base of the triangle, and $[z_0, z_3]$ is the *median* of the triangle.

Lemma 1. *Let M be a meniscus $M_\varepsilon(\theta)$ with base b and handle h, and let μ and s be the median and one of the two equal sides of the enclosing triangle of M. Then*

- $|h| = \varepsilon^2$ *and* $|b| \approx 2\varepsilon\sqrt{2}$.
- $\mathrm{Arc}(M) \approx 2\varepsilon\sqrt{2}$.
- $|\mu| \approx 2\varepsilon^2$.
- $|s| \approx \varepsilon\sqrt{2}$.

The approximate equalities mean that higher powers of ε are ignored.

Proof. Let the points z, z_0, z_1, z_2, z_3 be as above. The first claim follows from the definition of the meniscus. Let $x = \mathrm{Arc}(M)/2$.

To prove the second claim, note that $\cos x = |O, z_0|/|O, z_2| = |O, z_0| = 1 - \varepsilon^2$. Since the Taylor series for $\cos x$ is $1 - x^2/2 + \ldots$, we have $x \approx \varepsilon\sqrt{2}$ and $\mathrm{Arc}(M) \approx 2\varepsilon\sqrt{2} = O(\varepsilon)$.

Let $\delta = |z, z_3|$. We have $1-\varepsilon^2 = \cos x = |O, z_2|/|O, z_3| = 1/(1+\delta)$, so $1+\delta = 1/(1-\varepsilon^2) = 1 + \varepsilon^2 + \varepsilon^4 + \dots$ and $\delta \approx \varepsilon^2$. Since $\mu = [z_0, z_3]$, we have $|\mu| \approx 2\varepsilon^2$.

By definition of s, we have $|s|^2 = |\mu|^2 + (|b|/2)^2 \approx (|b|/2)^2$, so $|s| \approx |b|/2$. □

Recall that the trace distance $\mathrm{TD}(U, V)$ between unitary operators U, V (up to phase factors) of the Hilbert space $\mathbb{C}^2$ is $\sqrt{1 - |\mathrm{Tr}(UV^\dagger)|/2}$.

Lemma 2. *Let U be a unitary operator $\left(\begin{smallmatrix} u & -v^* \\ v & u^* \end{smallmatrix}\right)$ on $\mathbb{C}^2$, R be a z-rotation $R_z(\theta) = \left(\begin{smallmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{smallmatrix}\right)$, and M be the meniscus $M_\varepsilon(\theta)$. Then*

$$\mathrm{TD}(U, R) < \varepsilon \iff u_M \in M$$

where

$$u_M = \begin{cases} u & \text{if } \Re(ue^{i\theta/2}) \geq 0 \\ -u & \text{otherwise} \end{cases}$$

Proof. Note that $\mathrm{Tr}(UR^\dagger) = (ue^{i\theta/2}) + (ue^{i\theta/2})^* = 2\Re(ue^{i\theta/2})$.

$$\begin{aligned} \mathrm{TD}(U, R) < \varepsilon &\iff \sqrt{1 - |\mathrm{Tr}(UR^\dagger)|/2} < \varepsilon \\ &\iff 1 - |\mathrm{Tr}(UR^\dagger)|/2 < \varepsilon^2 \\ &\iff 2(1 - \varepsilon^2) < |\mathrm{Tr}(UR^\dagger)| \\ &\iff 1 - \varepsilon^2 < |\Re(ue^{i\theta/2})| \\ &\iff 1 - \varepsilon^2 < \Re(u_M e^{i\theta/2}) \\ &\iff 1 - \varepsilon^2 < \Re(u_M e^{i\theta/2}) \leq 1 \\ &\iff u_M e^{i\theta/2} \in M_\varepsilon(0) \\ &\iff u_M \in M_\varepsilon(\theta) \end{aligned}$$

The penultimate equivalence uses the fact that $|u| \leq 1$ which is true because u occurs in a unitary matrix. □

Corollary 3. *If $u \in M_\varepsilon(\theta)$ and there exists a complex number v satisfying the norm equation $|u|^2 + |v|^2 = 1$, then the matrix $U = \left(\begin{smallmatrix} u & -v^* \\ v & u^* \end{smallmatrix}\right)$ is at trace distance $< \varepsilon$ from the z-rotation $R = R_z(\theta)$.*

Proof. Since $u \in M_\varepsilon(\theta)$, we have $ue^{i\theta/2} \in U_0(\theta)$, $\Re(ue^{i\theta/2}) \geq 0$, and $u_M = u$. Now use Lemma 2. □

B. Pauli+V

We use the same letter for a linear operator on $\mathbb{C}^2$ and its matrix in the standard basis.

Unitary operators

$$V_1 = (I + 2iX)/\sqrt{5}, \quad V_2 = (I + 2iY)/\sqrt{5}, \quad V_3 = (I + 2iZ)/\sqrt{5}$$

generate a group V that is dense in the special unitary group SU(2) [Ref. 11]. Here I is the identity matrix, and X, Y, Z are the single-qubit Pauli matrices. We have

$$V_1^{-1} = (I - 2iX)/\sqrt{5}, \quad V_2^{-1} = (I - 2iY)/\sqrt{5}, \quad V_3^{-1} = (I - 2iZ)/\sqrt{5}.$$

The group V is in fact freely generated by V_1, V_2, V_3. This fact appears without proof in [Ref. 11]. For completeness, we prove it here; see Corollary 5 below.

In [Ref. 2], we explored a slightly larger group W generated by the three operators V_j and three Pauli operators X, Y, Z. Because of $Y(I + 2iX) = (I - 2iX)Y$ and similar relations, every product of operators in W easily reduces to a normal form

$$\begin{aligned} A_1 A_2 \cdots A_t B \quad \text{where} \quad & \sqrt{5} A_i \in \{1 \pm 2iX, 1 \pm 2iY, 1 \pm 2iZ\}, \\ & A_1 A_2 \cdots A_t \text{ is reduced, and} \\ & B \in \{\pm I, \pm X, \pm Y, \pm Z\}; \end{aligned} \tag{1}$$

in the process the number of V factors does not change but the number of Pauli factors becomes ≤ 1. The product $A_1 A_2 \cdots A_t$ is reduced in the sense that no A_{j+1} is the inverse of A_j.

By Theorem 1 in [Ref. 2], an SU(2) operator is in W if and only if it can be given in the form

$$\frac{1}{(\sqrt{5})^t} \begin{pmatrix} u & -v^* \\ v & u^* \end{pmatrix} \tag{2}$$

where u, v are Gaussian integers.

The next theorem will relate the exponent t in this formula to the number of factors in the normal form of an element of W. It will also provide similar information about the images in SO(3) of the Pauli+V matrices.

Recall how matrices in SU(2) act as rotations on three-dimensional Euclidean space. They act by conjugation on the 3-dimensional vector space of traceless Hermitian matrices

$$xX + yY + zZ = \begin{pmatrix} z & x - iy \\ x + iy & -z \end{pmatrix},$$

where X, Y, and Z are the Pauli matrices, and where we regard the real numbers x, y, z as coordinates in $\mathbb{R}^3$. Furthermore, this conjugation action preserves the Euclidean norm

$$x^2 + y^2 + z^2 = -\det \begin{pmatrix} z & x - iy \\ x + iy & -z \end{pmatrix},$$

so we get a homomorphism of SU(2) into the orthogonal group O(3). Because SU(2) is connected, the homomorphism actually maps into SO(3).

Under this homomorphism, the V-matrices correspond to rotations by $\arccos(-3/5)$ about the three coordinate axes, namely

$$R_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -\frac{3}{5} & \frac{4}{5} \\ 0 & -\frac{4}{5} & -\frac{3}{5} \end{pmatrix}, \quad R_2 = \begin{pmatrix} -\frac{3}{5} & 0 & -\frac{4}{5} \\ 0 & 1 & 0 \\ \frac{4}{5} & 0 & -\frac{3}{5} \end{pmatrix}, \quad R_3 = \begin{pmatrix} -\frac{3}{5} & \frac{4}{5} & 0 \\ -\frac{4}{5} & -\frac{3}{5} & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Theorem 4. *1. Any matrix obtained as a reduced product of t factors taken from $\{{R_1}^{\pm 1}, {R_2}^{\pm 1}, {R_3}^{\pm 1}\}$ has at least one entry which, when written as a fraction in lowest terms, has denominator 5^t.*

2. Any matrix obtained as a reduced product of t factors taken from $\{{V_1}^{\pm 1}, {V_2}^{\pm 1}, {V_3}^{\pm 1}\}$ has the form (2), and it cannot be written in that form with t replaced by a smaller exponent.

3. Any matrix in the normal form described above, with t factors taken from $\{{V_1}^{\pm 1}, {V_2}^{\pm 1}, {V_3}^{\pm 1}\}$ followed by one factor from $\{\pm I, \pm X, \pm Y, \pm Z\}$, has the form (2), and it cannot be written in that form with t replaced by a smaller exponent.

Proof. The proof of item (1) in the theorem is rather long and is therefore given in an appendix. We give here the easy deductions of items (2) and (3) from item (1).

To prove (2), consider any product $M = A_1 \cdots A_t$ of t factors, each of which is in $\{{V_1}^{\pm 1}, {V_2}^{\pm 1}, {V_3}^{\pm 1}\}$. Each factor is thus a matrix of Gaussian integers divided by $\sqrt{5}$, so the product M is a matrix of Gaussian integers dvided by $\sqrt{5^t}$. We need to show that no lesser power of $\sqrt{5}$ can serve as the denominator for M. So suppose, toward a contradiction, that M is a matrix of Gaussian integers divided by $\sqrt{5^r}$ with $r < t$. Then the the same is true of the conjugate transpose of M, which is also the inverse of M because M is unitary. Thus, when M acts by conjugation on the three-dimensional space of traceless Hermitian

matrices, the denominators are (at most) 5^r, namely a factor $\sqrt{5^r}$ from M and another factor $\sqrt{5^r}$ from M^{-1}. This means that the image of M in SO(3), which is given by this conjugation action on traceless Hermitian matrices, involves denominators only 5^r. But this element of SO(3) is obtained by multiplying the R matrices corresponding to the V matrices that produced M. So we would have a reduced product of t factors from $\{{R_1}^{\pm 1}, {R_2}^{\pm 1}, {R_3}^{\pm 1}\}$ with only 5^r in the denominator. This contradicts item (1).

Finally, for item (3), we must show that what we just proved about products of the form $A_1 \cdots A_t$ is also valid for $A_1 \cdots A_t B$ where $B \in \{\pm I, \pm X, \pm Y, \pm Z\}$. But this is easy, since the entries of B (and of B^{-1}, because $B^{-1} = B$) are Gaussian integers, so multiplication by B has no effect on the number of $\sqrt{5}$ factors needed in the denominator. □

Corollary 5. *1. The matrices R_1, R_2, and R_3 are free generators of the subgroup of* SO(3) *that they generate.*

2. The matrices V_1, V_2, and V_3 are free generators of the subgroup of SU(2) *that they generate.*

Proof. Both parts follow immediately from the corresponding parts of the theorem. The identity matrix cannot be represented by a nonempty reduced word in the given generators and their inverses, because it has no 5 or $\sqrt{5}$ in the denominator. □

Proposition 6. *The normal forms $A_1 \cdots A_t B$ all represent distinct matrices, so that every element of W has a unique normal form.*

Proof. Part (3) of Theorem 4 immediately implies that two normal forms with distinct lengths represent distinct matrices, for they have different powers of $\sqrt{5}$ in their simplest forms. So we need only consider normal forms of one weight t at a time. Fix t for the rest of this proof.

Combining Part (3) of Theorem 4 with Theorem 1 in [Ref. 2], we find that every matrix in W whose simplest form is (2) (with our fixed t) is represented by a normal form $A_1 \cdots A_t B$ (again with our fixed t). To show that this representation is unique, it suffices to show that the number of such matrices equals the number of such normal forms.

To count the relevant matrices (2), write $u = a + bi$ and $v = c + di$, where a, b, c, and d are ordinary integers, and observe that the matrix (2) is in SU(2) if and only if

$$a^2 + b^2 + c^2 + d^2 = |u|^2 + |v|^2 = 5^t.$$

Thus, the number of matrices in SU(2) of the form (2) is the number of representations of 5^t as a sum of four squares of integers. The number of such matrices for which this is the simplest form is then obtained by subtracting the number of such four-square representations in which all of a, b, c, and d are divisible by 5.

By Jacobi's four-square theorem, every positive odd integer n has $8\sum_{d|n} d$ representations as a sum of four squares of integers. In particular, for $n = 5^t$, there are

$$8(1 + 5 + \cdots + 5^t) = 8\frac{5^{t+1} - 1}{5 - 1} = 2(5^{t+1} - 1)$$

representations of 5^t as a sum $a^2+b^2+c^2+d^2$. As noted above, we must subtract the number of these representations in which all of a, b, c, and d are divisible by 5. Dividing these four integers by 5, we obtain the representations of 5^{t-2} as a sum of four squares, so the number to be subtracted is $2(5^{t-1} - 1)$. Therefore, the number of matrices whose simplest form is (2) is

$$2(5^{t+1} - 1) - 2(5^{t-1} - 1) = 2 \cdot 5^{t-1} \cdot (25 - 1) = 48 \cdot 5^{t-1}.$$

Now, let us count the number of normal forms $A_1 \cdots A_t B$. We have 6 choices for A_1 (namely any of the $V_i^{\pm 1}$), 5 choices for each subsequent A_j (namely any of the $V_i^{\pm 1}$ except the inverse of the immediately preceding A_{j-1}), and 8 choices for B. That makes $48 \cdot 5^{t-1}$ normal forms. Since this count agrees with the count of matrices above, the proof of the proposition is complete. □

The V-*count* of a W-operator U is t if U has a normal form $A_1 \cdots A_t B$. Every W-circuit implementing U contains at least t V-gates.

C. Diophantine approximations

We presume that the reader is familiar with continued fractions, and we use Khinchin's book [Ref. 8] as our reference on continued fractions.

Every rational number q has a unique continued-fraction representation $[a_0; a_1, a_2, \ldots, a_n]$ where all a_i are integers, $a_1, a_2, \ldots$ are positive and $a_n > 1$. Every irrational number has a unique continued-fraction representation $[a_0; a_1, a_2, \ldots]$ where all a_i are integers, and $a_1, a_2, \ldots$ are positive. If $[a_0; a_1, \ldots, a_k]$ is an initial segment of the continued-fraction representation of γ then the reduced fraction p_k/q_k represented by $[a_0; a_1, \ldots, a_k]$ is the k^{th} approximant (or convergent) of γ.

By a theorem of Dirichlet [Ref. 8, Theorem 25], for any real number γ and any integer $r \geq 1$, there exist relatively prime integers x and y such that

$$|\gamma y - x| < 1/r \quad \text{and} \quad 1 \leq y \leq r.$$

The proof of Theorem 25 in [Ref. 8] includes this claim: If p_k/q_k is the approximant of γ such that $q_k \leq r$ and either $r < q_{k+1}$ or else $\frac{p_k}{q_k} = \gamma$ then $|\gamma q_k - p_k| < 1/r$.

Lemma 7. *There is a polynomial-time algorithm that, given a rational g and integer $r \geq 1$, computes an approximant p_k/q_k of g such that $|gq_k - p_k| < 1/r$ and $1 \leq q_k \leq r$.*

Proof. The desired algorithm is recursive. Let $\gamma_0 = \gamma$, $a_0 = \lfloor \gamma_0 \rfloor$, $p_0 = a_0$, $q_0 = 1$, and suppose we computed already γ_j, a_j, p_j, q_j. If $\gamma_j - a_j < 1/r$, stop and output p_j/q_j. Otherwise let $\gamma_{j+1} = 1/(\gamma_j - a_j)$, $a_{j+1} = \lfloor \gamma_{j+1} \rfloor$ and

$$p_{j+1} = a_{j+1} p_j + p_{j-1}$$
$$q_{j+1} = a_{j+1} q_j + q_{j-1}$$

where $p_{-1} = 1$ and $q_{-1} = 0$.

To estimate the running time of the algorithm, use the fact that any $q_n \geq 2^{(n-1)/2}$ [Ref. 8, Theorem 12]. If the output is p_k/q_k, we have $2^{(k-1)/2} \leq q_k \leq r$ and $k \leq 1 + 2 \log_2 r$. □

Proviso 8. Every real number γ, used as input to an algorithm in the present paper, comes with an oracle that, given the unary notation for an integer $m \geq 0$, produces the part $\sum_{n=0}^{m} d_n/10^n$ of the decimal notation $\sum_{n=0}^{\infty} d_n/10^n$ for γ, where every d_n is an integer and $d_1, d_2, \ldots$ are in $\{0, 1, \ldots, 9\}$.

Lemma 9. *There is a polynomial-time algorithm that, given a real γ and integer $r \geq 1$, computes a reduced fraction x/y such that $|\gamma y - x| < 1/r$ and $y \leq 2r$.*

Proof. Use the oracle companion of γ to compute a rational g such that $|\gamma - g| \leq \frac{1}{2r^2}$. Use the algorithm of Lemma 7 to compute a fraction x/y such that $1 \leq y \leq 2r$ and $|\gamma y - x| \leq \frac{1}{2r}$. We have

$$|\gamma - \frac{x}{y}| \leq |\gamma - g| + |g - \frac{x}{y}| \leq \frac{1}{2r^2} + \frac{1}{2ry} \leq \frac{1}{ry}$$

and so $|\gamma y - x| \leq \frac{1}{r}$. □

D. Sums of squares

We recall some well-known facts related to the problem of representing a given (rational) integer $n \geq 0$ as a sum of two squares of integers. All variables will range over the integers. For brevity, we say that n is S2S if it is a sum of two squares.

Every prime number of the form $4m + 1$ is S2S. Given any such prime p, the Rabin-Shallit algorithm finds an S2S representation $x^2 + y^2$ of p in expected time $O(\log p)$ [Ref. 18, Theorem 11].

The S2S property is multiplicative. Indeed, if $m = a^2 + b^2 = |a + bi|^2$ and $n = x^2 + y^2 = |x + yi|^2$ then $mn = |(a + bi)(x + yi)|^2 = (ax - by)^2 + (ay + bx)^2$.

Lemma 10. *There is an algorithm that, given the representation of any number n as a product of powers of distinct primes, decides whether n is S2S and, if yes, produces an S2S representation $x^2 + y^2$ of n. The algorithm works in expected polynomial time.*

Proof. A number n is S2S if and only if every prime factor q of n of the form $4m+3$ has an even exponent in the representation of n as a product of powers of distinct primes [Ref. 5, Theorem 366]. This criterion allows you to decide whether n is S2S or not.

If n is S2S, use the Rabin-Shallit algorithm and the multiplicativity of S2S to find an S2S representation of n. We illustrate this part on an example. Suppose that $n = p_1 p_2^3 q^2$ where p_1, p_2 are primes of the form $4m + 1$ and q is a prime of the form $4m + 3$. Use the Rabin-Shallit algorithm to represent p_1, p_2 as sums of squares. Use the algorithm of the paragraph preceding the lemma, to represent $p_1 p_2$ as $c^2 + d^2$. Then $n = (cp_2q)^2 + (dp_2q)^2$. □

By the prime number theorem, the number $\pi(n)$ of primes $\leq n$ is asymptotically equal to $\frac{n}{\log n}$ where log means the natural logarithm. The number of primes of the form $4m+1$ that are $\leq n$ is asymptotically equal to $\frac{n}{2 \log n}$ and thus is $\Omega(\frac{n}{\log n})$. It follows that the fraction of S2S numbers $\leq n$ is $\Omega(\frac{1}{\log n})$.

Proposition 11. *There is an algorithm that, given a positive integer n of the form $4m + 1$ and a positive $\delta > 0$, works in time $O((\log n) \log \frac{1}{\delta})$ and returns an S2S representation of n or Nil. If n is prime then,with probability $> 1 - \delta$, the algorithm returns an S2S representation of n.*

Proof. Let A be the Rabin-Shallit algorithm for finding an S2S representation of a given prime number $n = 4m + 1$. A works in two stages. At stage 1, it solves the equation

$x^2 = -1$ in the field $\mathbb{Z}_n$. That solution is used at stage 2 to produce an S2S representation of n. Stage 2 is performed in linear time. Stage 1 is a while loop. At the j^{th} round of the loop, A randomly chooses a residue $b_j \mod n$ and computes, in linear time, the greatest common divisor $d_j(x) = x - a_j$ of polynomials $(x - b_j)^2 + 1$ and $x^{2m} - 1$. With probability $1/2$, the residue $a_j + b_j$ solves $x^2 = -1$; if this happens, call the round successful. The expected number of the rounds is 2. So A works in expected linear time.

Let A' be the modification of A that takes an additional input $\delta > 0$ and replaces the while loop with the following for-loop where $J = \lceil \log_2(1/\delta) \rceil$. For $j = 0$ to J do:

1. Perform one round of the A's while loop.

2. If the round is successful, go to stage 2.

3. If $j < J$ then increment j else stop and return Nil.

The probability that A' outputs Nil is $1/2^{J+1} < \delta$. The worst-case running time of A' is $O((\log n) \log \frac{1}{\delta})$.

The desired algorithm A'' is the modification of A' where the input integer $n = 4m + 1$ is not necessarily prime. A'' simulates A' on the given n and δ. If A' returns an alleged S2S representation $c^2 + d^2$ then A" checks whether the representation is genuine and returns the same S2S representation of n if it is genuine. In all other cases, A'' returns Nil. □

III. ADJUSTING A MENISCUS

We present an algorithm that, given a meniscus $M_\varepsilon(\theta)$ of the unit disk, constructs an operator $\tau \in \mathrm{SL}(2, \mathbb{Z})$ such that τM resides in a vertical band of width $O(\varepsilon^{3/2})$.

Call a complex number z *quasi-rational* if $\Re(z) = 0$ or $\Im(z)/\Re(z)$ is rational. Every nonzero quasi-rational z has a unique *reduced* presentation of the form $\mu(a + bi)$ where μ is real and positive and where a, b are mutually prime integers; if $\Re(z) \neq 0$ then b/a is the reduced form of the fraction $\Im(z)/\Re(z)$.

Observe that a vector r is orthogonal to a given quasi-rational vector $\mu(a + bi)$ if and only if r is quasi-rational of the form $\nu(b - ai)$.

Lemma 12. *Let q be a nonzero quasi-rational with reduced presentation $\mu(a + bi)$. For any nonzero complex number r orthogonal to q, there is $\tau \in \mathrm{SL}(2, \mathbb{Z})$ such that*

1. $\tau q = \mu i$ *and* $|\Im(\tau r)| < |\Re(\tau r)|$,

2. $|\tau q| = |q|/\sqrt{a^2+b^2}$,

3. $|\Re(\tau r)| = \sqrt{a^2+b^2}\,|r|$.

Proof of the lemma. 1. Use the extended Euclidean algorithm to find integers u, v with $ua + vb = 1$ and let $\tau_0 = \left(\begin{smallmatrix} b & -a \\ u & v \end{smallmatrix}\right)$. Then $\tau_0 q = \mu i$ and $\tau_0 r$ is some complex number $\alpha + \beta i$. Since q and r are non-collinear, so are τq and τr, and therefore $\alpha \neq 0$. If $|\beta| < |\alpha|$, the desired $\tau = \tau_0$. Otherwise $\tau = \tau'\tau_0$ where $\tau' = \left(\begin{smallmatrix} 1 & 0 \\ k & 1 \end{smallmatrix}\right)$ and $k = \lceil -\frac{\beta}{\alpha} \rceil = -\frac{\beta}{\alpha} + \delta$ for some δ such that $0 \leq \delta < 1$. We have $\tau q = \tau'(\mu i) = \mu i$, $\Re(\tau r) = \Re(\tau'(\alpha + \beta i)) = \alpha$, and $\Im(\tau r) = \Im(\tau'(\alpha + \beta i)) = k\alpha + \beta = \delta\alpha$, so that indeed $|\Im(\tau r)| < |\Re(\tau r)|$.

2. By the first part of Claim 1, $|\tau q| = |\mu| = |q|/\sqrt{a^2+b^2}$.

3. τ maps the rectangle with sides q and r to a parallelogram of the same area because $\tau \in \mathrm{SL}(2, \mathbb{Z})$. The parallelogram has a vertical side $\tau q = \mu i$ and side τr whose horizontal component is $\Re(\tau r)$, so its area is $|\mu| \cdot \Re(\tau r)$. The original rectangle had area $|q| \cdot |r| = |\mu|\sqrt{a^2+b^2}\,|r|$. Equating the two areas and cancelling $|\mu|$, we get the claim. □

Corollary 13. *Let M be a meniscus with a quasi-rational base vector q of reduced form $\mu(a+bi)$ and with handle h. There is $\tau \in \mathrm{SL}(2, \mathbb{Z})$ such that τq is vertical, $|\tau q| = |q|/\sqrt{a^2+b^2}$ and $|\Re(\tau h)| = \sqrt{a^2+b^2}\,|h|$.*

Originally, to achieve the goal of this section, we intended to show that, for every meniscus M of the unit circle, there exists a slightly bigger meniscus $L \supseteq M$ with a base vector q and a handle h and there exists an operator $\tau \in \mathrm{SL}(2, \mathbb{Z})$ such that τq is vertical and $|\Re(\tau h)| = O(\varepsilon^{3/2})$. The intent ran into difficulties with Diophantine approximations; see §II C in this connection. Fortunately there is another way.

Theorem 14. *Let M be a meniscus $M_\varepsilon(\theta)$ of the unit disk. There is $\tau \in \mathrm{SL}(2, \mathbb{Z})$ such that τM resides in a vertical band of width $O(\varepsilon^{3/2})$. Moreover, there is a polynomial-time algorithm that, given ε and θ, constructs the desired τ.*

Proof of Theorem 14. Since the isometry $\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right) \in \mathrm{SL}(2, \mathbb{Z})$ makes horizontal bands vertical, it suffices to prove the version of the theorem where "vertical" is replaced with "vertical or horizontal."

Let z_1, z_2 be the two corner points of M and $r = \alpha + \beta i$ be the base vector $z_2 - z_1$. We assume that $\alpha\beta \neq 0$; otherwise we have nothing to do. Without loss of generality, z_1 is the left of the two corner points of M, so that $\alpha > 0$. The enclosing triangle Z of M is formed by points z_1, z_2 and the intersection point, call it z_3, of the tangent lines to the unit circle at z_1 and z_2. Without loss of generality, $|\beta| \leq |\alpha|$; otherwise, instead of making Base(M) nearly vertical, we'll make it nearly horizontal (even though this may look a bit unnatural: if the base is closer to horizontal then we adjust it to become vertical, and if it's closer to vertical then we adjust it to become horizontal.)

We are going to construct an operator $\tau \in \mathrm{SL}(2, \mathbb{Z})$ such that the horizontal projections of all sides of the triangle τZ are $O(\varepsilon^{3/2})$. Let $\gamma = \beta/\alpha$ and $n = \lceil 1/\sqrt{\varepsilon}\,\rceil$. Apply the algorithm of Lemma 9 to construct a reduced fraction b/a such that $|\gamma a - b| < 1/n \leq \sqrt{\varepsilon}$ and $1 \leq a \leq 2n = 2\lceil 1/\sqrt{\varepsilon}\,\rceil$.

Since $|\gamma a - b| < 1$ and b is an integer, b and γa cannot have the opposite signs. Hence $b = 0$ or b has the sign of γa which is also the sign of γ and β. Recall that $\alpha \geq \beta$. We claim that $a \geq |b|$. Indeed,

$$|b| - a = |b| - a|\gamma| - a(1 - |\gamma|) \leq |b| - a|\gamma| = |b - \gamma a| < \sqrt{\varepsilon} < 1$$

and thus $|b| \leq a$.

Let q be the quasi-rational $\alpha + \frac{b}{a}\alpha i = \frac{\alpha}{a}(a + bi)$. Recall from §II A that $|r| = \Theta(\varepsilon)$. We have $r - q = (\beta - \frac{b}{a}\alpha)i$ and $|r - q| = |\frac{\alpha}{a}(\gamma a - b)| \leq \frac{|r|}{a}|\gamma a - b| = O(\varepsilon^{3/2}/a)$.

Consider meniscus $L = M_\delta(\eta) \supseteq M$ such that Base(L) is parallel to q and touches Base(M) at z_1 or z_2. We consider only the case $z_1 \in$ Base(L); the other case is similar. Let $y_1 = z_1$ and y_2 be the other end of Base(L). The enclosing triangle of L is formed by points y_1, y_2 and the intersection point, call it y_3, of the tangent lines to the unit circle at y_1 and y_2.

We have $\mathrm{Arc}(y_1, y_2) = \mathrm{Arc}(z_1, z_2) + \mathrm{Arc}(z_2, y_2) = O(\sqrt{\varepsilon}/a)$ and, by Lemma 1, $\delta = O(\sqrt{\varepsilon}/a)$. Let h be the handle of L. By Lemma 1, $|h| = \delta^2$. By Lemma 12, there is $\tau \in \mathrm{SL}(2, \mathbb{Z})$ that makes Base(L) vertical and such that $|\Re(\tau h)| = \Theta(a \cdot |h|) = \Theta(a\delta^2)$. Furthermore, a particular $\tau = \tau'\tau_0$ is constructed in the proof of the lemma; we are going to take advantage of that.

The horizontal projection of $[\tau z_1, \tau z_2]$ is $\Re(\tau r)$. We claim that the length of the horizontal projection is $O(\varepsilon^{3/2})$. Since τ' preserves the real part of any vector, it suffices to show that

$|\Re(\tau_0 r)| = O(\varepsilon^{3/2})$. We have

$$\Re\left[\left(\begin{smallmatrix} b & -a \\ u & v \end{smallmatrix}\right)\left(\begin{smallmatrix} \alpha & \beta \end{smallmatrix}\right)\right] = b\alpha - a\beta = \alpha(b - a\gamma).$$

So $|\Re(\tau_0 r)| \le |r| \cdot |a\gamma - b| = O(\varepsilon^{3/2})$. It remains to show that the length of the horizontal projection of $[\tau z_1, \tau z_3]$ is $O(\varepsilon^{3/2})$.

Let $z = z_3 - z_1$ and $y = y_3 - y_1$. By Lemma 1, $|z| \approx \varepsilon\sqrt{2}$ and $|y| \approx \delta\sqrt{2}$; so $\frac{|z|}{|y|} \approx \frac{\varepsilon}{\delta}$. Since the vectors y, z are collinear and operator τ is linear and projection operators are linear as well, $\frac{|\Re(\tau z)|}{|\Re(\tau y)|} \approx \frac{\varepsilon}{\delta}$. Since $\Re(\tau y) = \Re(\tau h)$, we have

$$|\Re(\tau z)| = \frac{\varepsilon}{\delta}\Re(\tau h) = \Theta\left(\frac{\varepsilon}{\delta} a\delta^2\right) = \Theta(a\varepsilon\delta) = O(\varepsilon^{3/2}).$$

It remains to estimate the running time of our algorithm. To this end we need only to estimate the time needed to compute integers a, b and then integers u, v such that $au + bv = 1$; the remaining work takes constant time. By Lemma 9, integers a, b are computed in polynomial time. The Euclidean algorithm that computes u, v runs in polynomial time as well. □

IV. DETERMINISTIC SEARCH

We explain the deterministic search, how it works and why. We do that essentially on the Pauli+V example. But, to simplify the exposition and minimize distractions, we abstract away some details. It would be easy to abstract away more details with the price of making the exposition a little more involved.

Consider a finite universal basis $\mathcal{B}$ for single-qubit gates. $\mathcal{B}$ may contain gates considered negligible gates; in such a case the depth of a $\mathcal{B}$-circuit is the number of non-negligible gates in the circuit. In the Pauli+V case, the Pauli gates may be considered negligible, because they are relatively cheap to implement and because at most one Pauli gate occurs in the normal form (1).

Assume that $\mathcal{B}$ comes with a partial function Level(u) that assigns nonnegative integers to some complex numbers and with an equation $\mathrm{NE}_t(u, v)$ such that the following constraints C1-C4′ hold.

C1. If u, v are of levels $\le t$ and $\mathrm{NE}_t(u, v)$ holds, then the matrix $\left(\begin{smallmatrix} u & -v^* \\ v & u^* \end{smallmatrix}\right)$ is unitary and exactly realizable by a $\mathcal{B}$ circuit of depth $\le t$.

The "NE" in "$\mathrm{NE}_t(u,v)$" alludes to the fact that the condition is typically expressed as a norm equation on u and v, with parameter t. Whether $\mathrm{NE}_t(u,v)$ is a pure norm equation or not, we will refer to it as the *norm equation*. The reader may be interested how the Pauli+V fits the general scheme, in particular what are the levels and norm equation in the Pauli+V scheme; all these questions will be addressed in the next section.

Recall that we are seeking to approximate a given z-rotation $R_z(\theta) = \left(\begin{smallmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{smallmatrix}\right)$ to a given precision ε. Let $M = M_\varepsilon(\theta)$.

If $u \in M$ and $\mathrm{Level}(u) \le t$, call u a *candidate of level* $\le t$. By Corollary 3, if u is a candidate of level $\le t$ and v is a complex number of level $\le t$ such that $\mathrm{NE}_t(u,v)$ holds, then the matrix $U = \left(\begin{smallmatrix} u & -v^* \\ v & u^* \end{smallmatrix}\right)$ is at a distance $< \varepsilon$ from $R_z(\theta)$. Accordingly, call a candidate u of level $\le t$ a *winning candidate of level* $\le t$ (or simply a *winning candidate* if t is clear from the context) if there is a complex number v of level $\le t$ such that $\mathrm{NE}_t(u,v)$ holds.

Let C_t be the set of all candidates of levels $\le t$, and let W_t be the subset of all winning candidates of levels $\le t$. Assume that the following constraints C2–C4 hold. By default, log means natural logarithm.

C2. There is an efficient algorithm that, given t, enumerates the candidates of level t.

C3. There exist a real $a > 1$ and an integer t_0 such that $|C_{t_0}| > 1$, and $|C_t| \le 1$ for all $t < t_0$, and $|C_t| \ge a^{t-t_0}$ for all $t > t_0$ such that $t - t_0$ is even. Here t_0 depends on ε and θ while a depends on neither.

C4. $|W_t| = \Omega(|C_t|/t)$ as $t \to \infty$, uniformly with respect to ε.

The requirement that $t - t_0$ be even in C3 reflects a peculiarity of the Pauli+V case.

Lemma 15. *$|W_t| \ge k$ for some $t = t_0 + c \log t_0$ where $c \ge 0$ depends on k and ε but not θ.*

Proof. Let a, t_0 be as in C3. By C4, there exist a real $b > 0$ and an integer $t_1 \ge 1$, independent of k, ε or θ, such that $|W_t| \ge b|C_t|/t$ for all $t \ge t_1$.

Let $t \ge t_0$, $t \ge t_1$ and $t - t_0$ be even. Then $|W_t| \ge b|C_t|/t \ge ba^{t-t_0}/t$, and so $|W_t| \ge k$ if $a^{t-t_0} \ge kt/b$. If $t = t_0 + x \log t_0$ then

$$|W_t| \ge k \text{ if } (a^{\log t_0})^x \ge \frac{kt_0}{b} + \frac{k \log t_0}{b} x. \tag{3}$$

The exponential function $(a^{\log t_0})^x$ of x quickly outgrows the linear function $\frac{kt_0}{b} + \frac{k \log t_0}{b} x$ of x. The desired t is $t_0 + c \log t_0$ where c is the least nonnegative real such that $t_0 + c \log t_0 \ge t_1$, $x \log t_0$ is an even integer and the premise of the implication (3) holds. □

Our goal is to (efficiently) find a winning candidate, preferably of low level. Our ability to do this depends on our ability to tell whether a given candidate is winning, and in this connection we consider two scenarios.

Scenario 1 We have a deterministic decision procedure that, given an integer $t \geq 0$ and a candidate $u \in C_t$, decides in polynomial time whether $u \in W_t$. Then the following obvious deterministic search finds a winning candidate of the minimal possible level. Explore candidates of levels 0, then candidates of level 1, etc. until a winning candidate of some level is found. The efficiency of such a deterministic search crucially depends on the efficiency of the decision procedure.

Scenario 2 We have a randomizing procedure that, given an integer $t \geq 0$ and a candidate $u \in C_t$, decides in expected polynomial time whether u belongs to a subset W'_t of W_t subject to the following constraint.

C4′ $|W'_t| = \Omega(|C_t|/t)$ as $t \to \infty$, uniformly with respect to ε.

Then the following randomizing search finds a candidate in W'_t of the minimal possible level. Explore candidates of levels 0, then candidates of level 1, etc. until, for some t, a member of W'_t is found.

Lemma 16. *$|W'_t| \geq k$ for some $t = t_0 + c \log t_0$ where $c \geq 0$ depends on k and ε but not θ.*

Proof. Just replace the reference to C4 with a reference to C4′ in the proof of Lemma 15. □

V. OPTIMAL PAULI+V CIRCUITS

A. Pauli+V candidates

We specialize §IV to the Pauli+V case. All the assumptions made in §IV need to be justified.

Define a complex number u to be of level $\leq t$ if $\sqrt{5^s}u$ is a Gaussian integer for some nonnegative integer $s \leq t$. The norm equation $\mathrm{NE}_t(u, v)$ has a particularly simple form in the Pauli+V case: $|u|^2 + |v^2| = 1$. If $\sqrt{5^t}u, \sqrt{5^t}v$ are Gaussian integers $a + bi, c + di$ respectively then the norm equation becomes $a^2 + b^2 + c^2 + d^2 = 5^t$. By Theorem 1 in [Ref. 2], mentioned in § II B, constraint C1 holds.

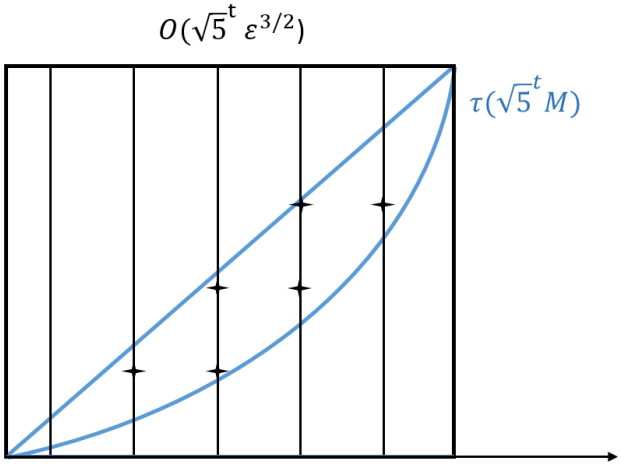

FIG. 1: Enumeration of Gaussian integers in $\sqrt{5^t}\tau M$.

Toward verifying C2, construct a linear transformation τ as in Theorem 14. The transformation τ maps straight lines into straight lines and ellipses into ellipses; it preserves areas and convexity. The elliptical meniscus τM is enclosed in a vertical band of width $O(\varepsilon^{3/2})$.

The inflated elliptical meniscus $\sqrt{5^t}\tau M$ is enclosed in a vertical band that projects onto the real segment $[l_t, r_t]$ with $r_t - l_t = O(\sqrt{5^t}\varepsilon^{3/2})$. $\sqrt{5^t}\tau M$ is bounded by segments

$$a + if(a) : \ l_t \le a \le r_t$$
$$a + ig(a) : \ l_t \le a \le r_t$$

of a straight line and of an ellipse respectively. To simplify the exposition, we consider only the case where the straight line segment is above the ellipsis segment.

Each Gaussian integer in $\sqrt{5^t}\tau M$ belongs to a vertical segment $[n+ig(n), n+if(n)]$ where n is an integer in the segment $[l_t, r_t]$; see Figure 1. This allows us to enumerate efficiently all Gaussian integers $a+bi$ in $\sqrt{5^t}\tau M$ and thus to enumerate efficiently all candidates $\frac{\tau^{-1}(a+bi)}{\sqrt{5^t}}$ of levels $\le t$. Constraint C2 holds.

Constraint C3 follows from the following claim based on an observation in [Ref. 20].

Claim 17. *If $|C_t| \ge 2$ then for any integer $k \ge 0$ we have $|C_{t+2k}| \ge 1 + 5^k$.*

Proof. Suppose that z_1, z_2 are Gaussian integers and candidates $\frac{z_1}{\sqrt{5^t}}, \frac{z_2}{\sqrt{5^t}}$ of levels $\le t$ belong to M. For any k, these candidates are also of levels $\le t+2k$. Since M is convex, it contains

also the following intermediate candidates of levels $\leq t + 2k$:

$$\frac{5^t - j}{5^k} z_1 + \frac{j}{5^k} z_2 \quad \text{where } j = 1, ..., 5^k - 1.$$

□

Note that a candidate $u = \sqrt{5^{-t}}(a + bi)$ of level $\leq t$ is a winning candidate of level $\leq t$ if and only if $5^t - a^2 - b^2$ is a sum $c^2 + d^2$ of two squares. Here a, b, c, d are all integers, Constraint C4 follows from the following number-theoretic conjecture of [Ref. 2].

Conjecture 1. *Let A be the area of the meniscus $\sqrt{5^t} M_\varepsilon(\theta)$ of the disk of complex numbers of norm $\leq \sqrt{5^t}$, and let S be the set of Gaussian integers $a + bi$ in $\sqrt{5^t} M_\varepsilon(\theta)$ such that $5^t - a^2 - b^2$ is a sum of two squares (of rational integers). Then $|S| = \Omega(A/t)$.*

Actually, instead of Ω, one finds Θ in [Ref. 2], but only the lower bound is used there and also here.

Define W'_t to be the set of members $u = \sqrt{5^{-t}}(a + bi)$ of W_t such that $5^t - a^2 - b^2$ is a prime of the form $4m + 1$. Every such prime is a sum of two squares; see §II D.

To justify C4′, we need an additional number-theoretic conjecture.

Conjecture 2. *Let A, S be as in Conjecture 1, and let S' consists of numbers $a + bi \in S$ such that $5^t - a^2 - b^2$ is a prime of the form $4m + 1$. Then $|S'| = \Omega(A/t)$.*

Both conjectures were found credible by the experts in analytic number theory that we consulted. The conjectures also are supported by experimentation. The intuition behind the conjectures is that there is no correlation between sets S, S' one the one side and prime numbers on the other. As far as sets S and S' are concerned, the distribution of prime numbers could be random. "It is evident that the primes are randomly distributed but, unfortunately, we don't know what 'random' means" quipped the number theorist Bob Vaughan in 1990 [Ref. 13].

B. The first circuit-synthesis algorithm

Our first circuit-synthesis algorithm is presented in Figure 2 where P is a procedure that takes a positive integer n as input and returns a complex number $c + di$ with $c^2 + d^2 = n$

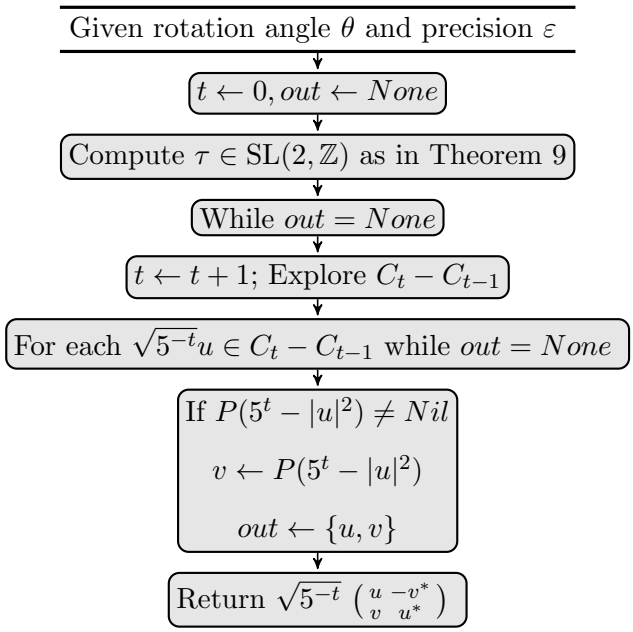
Given rotation angle θ and precision ε
$t \leftarrow 0, out \leftarrow None$
Compute $\tau \in \mathrm{SL}(2,\mathbb{Z})$ as in Theorem 9
While $out = None$
$t \leftarrow t+1$; Explore $C_t - C_{t-1}$
For each $\sqrt{5^{-t}}u \in C_t - C_{t-1}$ while $out = None$
If $P(5^t - |u|^2) \neq Nil$
$v \leftarrow P(5^t - |u|^2)$
$out \leftarrow \{u, v\}$
Return $\sqrt{5^{-t}} \begin{pmatrix} u & -v^* \\ v & u^* \end{pmatrix}$

FIG. 2: First synthesis algorithm

or returns Nil. We give 2 variants of the synthesis algorithm that correspond to the two scenarios of §IV. The two variants differ only in their versions P_1, P_2 of the procedure P.

Procedure P_1

This (and only this) version of P presumes the availability of a quantum computer.

1. Use Shor's algorithm [Ref. 21] to factor the given number n. Shor's algorithm works in polynomial time on a quantum computer.

2. Return Nil if some prime factor of n of the form $4m+3$ has an odd exponent in the representation of n as a product of powers of distinct primes.

3. Use the algorithm of Lemma 10 to find a representation $n = c^2 + d^2$ of n as a sum of two squares; return $c + di$.

Procedure P_2

1. Check whether n has the form $4m+1$, and return Nil if not.

2. Use the Lenstra-Pomerance version of the AKS primality test that runs in time $(\log n)^6 \cdot (2 + \log\log n)^c$ for some real constant c [Ref. 10]. (Alternatively use Miller's primality test that runs in time $O(\log n)^4$ but assumes the Extended Riemann Hypothesis [Ref. 14].) Return Nil if n is composite.

3. Use the Rabin-Shallit algorithm to find a representation $n = c^2 + d^2$ of n; return $c + di$.

The first variant of the synthesis algorithm runs in expected polynomial time and produces an approximating circuit of the least possible depth.

Theorem 18. *Given a target angle θ and real $\varepsilon > 0$, the second variant of the synthesis algorithm works in expected polynomial time and constructs a Pauli+V circuit of depth at most $3\log_5 \frac{1}{\varepsilon} + O(\log\log \frac{1}{\varepsilon})$ that approximates the axial rotation $R_z(\theta)$ within ε.*

Proof.
Correctness It should be obvious at this point that the algorithm produces a circuit that approximates $R_z(\theta)$ within ε.

Circuit depth By Lemma 16, the synthesis algorithm produces a circuit of depth at most $t_0 + O(\log t_0)$ where t_0 is the least index t such that $|C_t| \geq 2$. It suffices to prove the following lemma.

Lemma 19. $t_0 \leq 3\log_5(1/\varepsilon)$ *for sufficiently small positive ε.*

Proof of Lemma 19. Let $M = M_\varepsilon(\theta)$. By Lemma 1, $\mathrm{Area}(M) \geq c\varepsilon^3$ where $c \approx 2\sqrt{2}$. Let $s = 3\log_5(1/\varepsilon)$. Then

$$\mathrm{Area}(\sqrt{5^s}M) \geq 5^s \cdot c\varepsilon^3 = (1/\varepsilon)^3 c\varepsilon^3 = c > 2.$$

The number $N(t)$ of Gaussian integers in $\sqrt{5^t}M$ is asymptotically equal to $\mathrm{Area}(\sqrt{5^t}M)$ as $t \to \infty$ [Ref. 7]. So, for sufficiently small positive ε, $N(s) > 2$ and $\sqrt{5^s}M$ contains at least 2 Gaussian integers. □

Running time By Lemma 16, W'_t contains at least two elements for some t of the form $t_0 + b\log t_0$ where $b \geq 0$. Accordingly, the synthesis algorithm needs to explore only candidates of levels $\leq t$. Each such candidate has the form $\frac{u}{\sqrt{5^t}}$ where u is a Gaussian integer in $\sqrt{5^t}M$. The number of such Gaussian integers is asymptotically equal to $\mathrm{Area}(\sqrt{5^t}M) = 5^t\mathrm{Area}(M) \leq 5^t 4\sqrt{2}\varepsilon^3$ where the inequality follows from Lemma 1. By Lemma 19, $t = 3\log_5(1/\varepsilon) + d\log_5\log(1/\varepsilon)$ for some $d \geq 0$, so that the algorithm needs to explore at most

$$(1/\varepsilon)^3 \log^d(1/\varepsilon) 4\sqrt{2}\varepsilon^3 = \Theta(\log^d(1/\varepsilon))$$

candidates, that is only a polynomial number of candidates.

In the previous subsection, we explained how to enumerate all the candidates in C_t. To this end we have to walk through all the vertical segments $[n + ig(n), n + if(n)]$ where $l_t \leq n \leq r_t$. There are only polynomially many of those vertical segments, namely

$$\sqrt{5^t}\varepsilon^{3/2} = (1/\varepsilon)^{3/2} \log^{d/2}(1/\varepsilon)\varepsilon^{3/2} = \log^{d/2}(1/\varepsilon). \tag{4}$$

Each vertical segment contains only polynomially many Gaussian integers, and the exploration of one candidate involves a procedure P that runs in expected polynomial time. The whole algorithm runs in expected polynomial time □

Remark 20. The reader may wonder whether the transformation τ was necessary. It was. The original M is enclosed in a vertical band of width that may be of the order of ε. Replacing $\varepsilon^{3/2}$ with ε gives us exponentially many vertical segments.

C. The second circuit-synthesis algorithm

The only essential difference of our second synthesis algorithm SA2 from the first one is its version of procedure P whose input includes a positive real δ in addition to an integer. But this influences the forms of input and output of SA2. The input of SA2 comprises three components: a target angle θ, a precision $\varepsilon > 0$ and an error tolerance $\delta > 0$. SA2 returns either an approximation to the rotation R_z or Nil. The probability of returning Nil is at most δ. SA2 is presented in Figure 3.

But, before we describe the new version of P, let's address Rabin's primality test [Ref. 16]. It is an efficient polynomial-time primality test with a parameter k. If n is prime then the test result is always correct. For a composite n the test may declare n to be prime, but the probability of such error is tiny. "The algorithm produces and employs certain random integers $0 < b_1, \ldots, b_k < n \ldots$" writes Rabin in [Ref. 16], "the probability that a composite number will be erroneously asserted to be prime is smaller than $\frac{1}{2^{2k}}$. If, say, $k = 50$, then the probability of error is at most $\frac{1}{2^{100}}$."

Now we are ready to describe the new version of procedure P.

1. Check whether n has the form $4m + 1$, and return Nil if not.

2. Use Rabin's primality test with $k \geq \frac{1}{2} \log \frac{1}{\delta}$. Return Nil if n is found to be composite.

3. Apply the algorithm of Proposition 11 to n and δ.

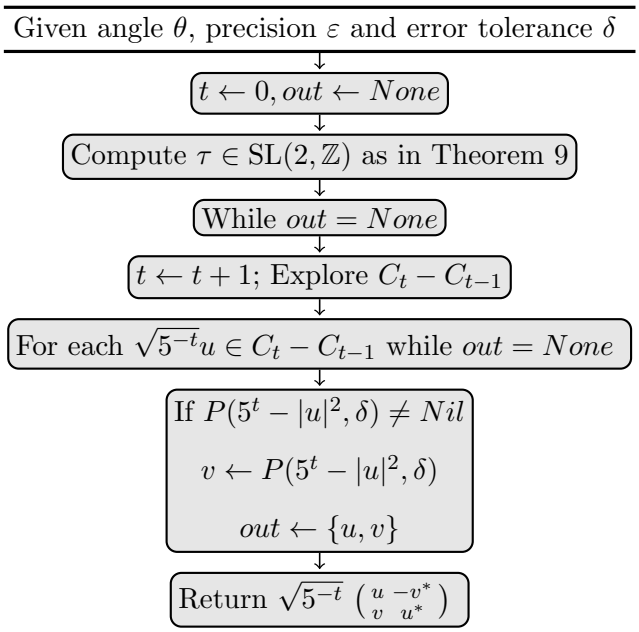
Given angle θ, precision ε and error tolerance δ
$t \leftarrow 0, out \leftarrow None$
Compute $\tau \in \mathrm{SL}(2,\mathbb{Z})$ as in Theorem 9
While $out = None$
$t \leftarrow t+1$; Explore $C_t - C_{t-1}$
For each $\sqrt{5^{-t}}u \in C_t - C_{t-1}$ while $out = None$
If $P(5^t - |u|^2, \delta) \neq Nil$
$v \leftarrow P(5^t - |u|^2, \delta)$
$out \leftarrow \{u, v\}$
Return $\sqrt{5^{-t}} \begin{pmatrix} u & -v^* \\ v & u^* \end{pmatrix}$

FIG. 3: Second synthesis algorithm

Corollary 21. *The new procedure P returns an S2S representation of n or Nil. The probability of returning Nil is at most δ.*

Theorem 22. *Given a target angle θ and real $\varepsilon, \delta > 0$, the second synthesis algorithm SA2 works in polynomial time and returns Nil or constructs a Pauli+V circuit of depth at most $3\log_5 \frac{1}{\varepsilon} + O(\log\log \frac{1}{\varepsilon})$ that approximates the axial rotation $R_z(\theta)$ within ε. The probability of returning Nil is at most δ.*

Proof. If the primality test works as intended, SA2 works essentially like the first one. If the primality test errs, which happens with probability at most δ, SA2 returns Nil or constructs a Pauli+V circuit of depth at most $3\log_5(1/\varepsilon) + O(\log\log \frac{1}{\varepsilon})$ that approximates the axial rotation $R_z(\theta)$ within ε. □

Appendix A: Three Free Rotations

In this appendix, we prove item (1) of Theorem 4. We shall be concerned with reduced words w built from the formal symbols (letters) l_1, l_2, l_3, and their (formal) inverses. (As in the theorem, "reduced" means that no l_i occurs next to its own inverse in w.) We write $w[R]$ for the product obtained by replacing each formal symbol l_i by the corresponding matrix R_i

exhibited just before the statement of Theorem 4, and replacing ${l_i}^{-1}$ by the inverse matrix (also the transpose, as the matrices are orthogonal). We must show that, if a reduced word w has length t, then the corresponding matrix $w[R]$ contains at least one entry which, when written as a fraction in lowest terms, has denominator 5^t.

We shall prove this result by induction on the length t of the word. It is clearly true for $t = 0$ and $t = 1$. For the induction step, we begin with some preliminary considerations to simplify the problem.

Each of the matrices R_i and each of their inverses can be written as $1/5$ times an integer matrix S_i, namely

$$S_1 = \begin{pmatrix} 5 & 0 & 0 \\ 0 & -3 & 4 \\ 0 & -4 & -3 \end{pmatrix}, \quad S_2 = \begin{pmatrix} -3 & 0 & -4 \\ 0 & 5 & 0 \\ 4 & 0 & -3 \end{pmatrix}, \quad S_3 = \begin{pmatrix} -3 & 4 & 0 \\ -4 & -3 & 0 \\ 0 & 0 & 5 \end{pmatrix}.$$

and $S_1^\top, S_2^\top, S_3^\top$. Note that we are factoring $1/5$ out of each R_i and each of the inverses $R_i^{-1} = R_i^\top$, so the remaining factors are the matrices S_i and $S_i^\top$, not S_i^{-1} (which differs from $S_i^\top$ by a factor 25).

Thus, if a reduced word w has length t, then $w[R] = (1/5^t)w[S]$, where $w[S]$ is obtained from w by replacing the letters l_i and their inverses l_i^{-1} in w by the matrices S_i and their transposes $S_i^\top$. What we must prove is that, in such a product $w[S]$ of S matrices, we never have all the entries divisible by 5; then, in any entry that is not divisible by 5, the overall factor of $1/5^t$ provides the denominator required for our result.

We have thus reduced our task to showing that, if we take a reduced word w and substitute for each letter l_i the corresponding S_i and for each ${l_i}^{-1}$ the transpose $S_i^\top$ of the corresponding S_i, then the product matrix $w[S]$ will not have all its entries divisible by 5.

We can reduce the task further. Since we are interested only in divisibility by 5, we can reduce all entries in the S matrices, and their transposes, and their products, modulo 5. That is, we can perform the whole calculation using matrices with entries in $\mathbb{Z}/5$, namely the matrices

$$T_1 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 2 & 4 \\ 0 & 1 & 2 \end{pmatrix}, \quad T_2 = \begin{pmatrix} 2 & 0 & 1 \\ 0 & 0 & 0 \\ 4 & 0 & 2 \end{pmatrix}, \quad T_3 = \begin{pmatrix} 2 & 4 & 0 \\ 1 & 2 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

obtained by reducing the S matrices modulo 5. As with the S matrices, we use the notation $w[T]$ for the result of taking a word w, replacing each l_i and l_i^{-1} with T_i and the transpose

$T_i^\top$, and multiplying the resulting matrices. Thus, $w[T]$ is a 3×3 matrix over the 5-element field $\mathbb{Z}/5$.

Our goal, that the product $w[S]$ of S matrices obtained from a reduced word w does not have all entries divisible by 5, can now be reformulated as saying that the product $w[T]$ of T matrices obtained from a reduced word is not the zero matrix.

Not only are the T matrices singular (because each has a column of zeros) but they have rank only 1, because, in each of them, the two non-zero columns are proportional. The same goes for the transposes of these matrices (either by similar inspection or because transposing a matrix doesn't change its rank). Let us write L_i for the one-dimensional subspace of $(\mathbb{Z}/5)^3$ that is the range of T_i and L_i' for the range of the transpose $T_i^\top$. Thus, L_i (resp. L_i') is generated as a vector space by either of the non-zero columns (resp. rows) of T_i.

With these preparations, we are ready for the induction step in the proof. Suppose the claim is true for reduced words of a certain length $t \geq 1$, and suppose we are given a reduced word of length $t+1$, say QW, where Q is the first letter of our given word and w is all the rest, i.e., a word of length t. The first letter of w (which exists as $t \geq 1$) is either some l_i or some l_i^{-1}.

Let us consider first the case that the first letter of w is l_i. (The case of an inverse l_i^{-1} will be similar.) By induction hypothesis, the matrix $w[T]$ is not the zero matrix. Its range is included in the range L_i of T_i (because the range of any matrix product JK is included in the range of the left factor J) and must therefore be all of L_i (since L_i has dimension only 1). To show that QW also corresponds to a non-zero matrix, it therefore suffices to show that $Q[T]$ does not annihilate L_i. (As Q is a single letter or inverse, $Q[T]$ is a single T matrix or transpose.) Here it is important that Q is not l_i^{-1}, because QW is a reduced word. So we need only check that the range L_i of T_i is not annihilated by any of the T_j's or their transposes, except for $T_i^\top$, and this is just a matter of inspection (the many 0's in the matrices make the computations trivial).

In the case where the first letter of w is some l_i^{-1}, we similarly reduce the problem to showing that the range L_i' of $T_i^\top$ is not annihilated by any of the T_j's or their transposes except for T_i. This again can be done by inspection, thereby completing the induction and thus the proof of Theorem 4(1).

REFERENCES

[1]M. Agrawal, N. Kayal and N. Saxena, "PRIMES is in P," Ann. Math. 160, 781–793. (2004)

[2]A. Bocharov, Y. Gurevich and K. Svore, "Efficient decomposition of single-qubit gates into V basis circuits," Physical Review A 88, 012313 (2013)

[3]A. Bocharov, M. Roetteler, and K. Svore, "Efficient synthesis of probabilistic quantum circuits with fallback," arXiv 1409.3552 (2014)

[4]M. Grötschel, L. Lovász and A. Schrijver, "Geometric algorithms and combinatorial optimization," 2nd corrected edition, Springer 1994

[5]G. Hardy and E. Wright, "An Introduction to the Theory of Numbers," Clarendon Press, 1979

[6]A. W. Harrow, B. Recht and I. L. Chuang, "Efficient discrete approximations of quantum gates," J. Math. Phys. 43:9, 4445 (2002)

[7]M. N. Huxley, "Area, Lattice points and Exponential Sums," Clarendon Press, 1996

[8]A. Y. Khinchin, "Continued Fractions," The University of Chicago Press, 1964

[9]V. Kliuchnikov, A. Bocharov and K. Svore, "Asymptotically optimal topological quantum compiling, arXiv 1310.4150 (2013)

[10]H. W. Lenstra, Jr. and C. Pomerance, "Primality testing with Gaussian periods," preprint, 2011. `http://www.math.dartmouth.edu/~carlp/aks041411.pdf`

[11]A. Lubotsky, R. Phillips and P. Sarnak, "Hecke operators and distributing points on S^2, I," Comm. Pure and Appl. Math. 34, 149–186 (1986)

[12]A. Lubotsky, R. Phillips and P. Sarnak, "Hecke operators and distributing points on S^2, II," Comm. Pure and Appl. Math. 40, 401–420 (1987)

[13]Mathematics Stack Exchange, #421353, "Are primes randomly distributed?" (2013), `http://math.stackexchange.com/questions/421353/`

[14]G. L. Miller, "Riemann's hypothesis and tests for primality," Computer and System Sciences 13, 300–317 (1976)

[15]M. A. Nielsen and I. Chuang, "Quantum Computation and Quantum Information," Cambridge University Press, 2000.

[16]M. O. Rabin, "Probabilistic algorithm for testing primality," Journal of Number Theory 12, 128–138 (1980)

[17]M. O. Rabin, "Probabilistic algorithms in finite fields," SIAM J. Comput. 9:2, 273–280

(1980)

[18]M. O. Rabin and J. O. Shallit, "Randomized algorithms in number theory," Comm. Pure and Appl. Math. 39, 239–256 (1986)

[19]N. J. Ross, "Optimal ancilla-free Pauli+V approximation of z-rotations," arXiv 1409.4355 (2014)

[20]N. J. Ross and P. Selinger, "Optimal ancilla-free Clifford+T approximation of z-rotations," arXiv 1403.2975 (2014)

[21]P. W. Shor, "Polynomial-time algorithm for prime factorization and discrete logarithms on a quantum computer," SIAM Journal on Computing 26, 1484–1509 (1997)

[22]P. Selinger, "Efficient Clifford+T approximation of single-qubit operators," arXive 1212.6253 (2012)